


Chapter 2 Application Layer



A note on the use of these ppt slides:

We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a lot of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) in substantially unaltered form, that you mention their source (after all, we'd like people to use our book!)
- If you post any slides in substantially unaltered form on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

All material copyright 1996-2010
J.F. Kurose and K.W. Ross, All Rights Reserved

*Computer Networking:
A Top Down Approach,
5th edition.
Jim Kurose, Keith Ross
Addison-Wesley, April
2009.*

Application 2-1

Chapter 2: Application layer

- | | |
|--|---------------------------------|
| 2.1 Principles of network applications | 2.6 P2P applications |
| 2.2 Web and HTTP | 2.7 Socket programming with TCP |
| 2.3 FTP | 2.8 Socket programming with UDP |
| 2.4 Electronic Mail <ul style="list-style-type: none">▪ SMTP, POP3, IMAP | |
| 2.5 DNS | |

Application 2-2

Chapter 2: Application Layer

Our goals:

- ❖ conceptual, implementation aspects of network application protocols
 - transport-layer service models
 - client-server paradigm
 - peer-to-peer paradigm
- ❖ learn about protocols by examining popular application-level protocols
 - HTTP
 - FTP
 - SMTP / POP3 / IMAP
 - DNS
- ❖ programming network applications
 - socket API

Application 2-3

Some network apps

- ❖ e-mail
- ❖ web
- ❖ instant messaging
- ❖ remote login
- ❖ P2P file sharing
- ❖ multi-user network games
- ❖ streaming stored video (YouTube)
- ❖ voice over IP
- ❖ real-time video conferencing
- ❖ cloud computing
- ❖ ...
- ❖ ...
- ❖ ...

Application 2-4


Creating a network app

write programs that

- run on (different) *end systems*
- communicate over network
- e.g., web server software communicates with browser software

No need to write software for network-core devices

- network-core devices do not run user applications
- applications on end systems allows for rapid app development, propagation



Application 2-5

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P applications
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP


Application 2-6

Application architectures

- ❖ client-server
- ❖ peer-to-peer (P2P)
- ❖ hybrid of client-server and P2P

Application 2-7

Client-server architecture



server:

- always-on host
- permanent IP address
- server farms for scaling

clients:


- communicate with server
- may be intermittently connected
- may have dynamic IP addresses
- do not communicate directly with each other

Application 2-8

Pure P2P architecture

- ❖ no always-on server
- ❖ arbitrary end systems directly communicate
- ❖ peers are intermittently connected and change IP addresses

highly scalable but
difficult to manage



Application 2-9

Hybrid of client-server and P2P

Skype

- voice-over-IP P2P application
- centralized server: finding address of remote party:
- client-client connection: direct (not through server)

Instant messaging

- chatting between two users is P2P
- centralized service: client presence detection/location
 - user registers its IP address with central server when it comes online
 - user contacts central server to find IP addresses of buddies

Application 2-10

Processes communicating

process: program running within a host.

- ❖ within same host, two processes communicate using **inter-process communication** (defined by OS).
- ❖ processes in different hosts communicate by exchanging **messages**

client process: process that initiates communication

server process: process that waits to be contacted

- ❖ aside: applications with P2P architectures have client processes & server processes

Application 2-11


Sockets

- ❖ process sends/receives messages to/from its **socket**

- ❖ socket analogous to door

- sending process shoves message out door
- sending process relies on transport infrastructure on other side of door which brings message to socket at receiving process

- ❖ API: (1) choice of transport protocol; (2) ability to fix a few parameters (lots more on this later)



Application 2-12

Addressing processes

- ❖ to receive messages, process must have *identifier*
- ❖ host device has unique 32-bit IP address
- ❖ *Q:* does IP address of host on which process runs suffice for identifying the process?

Application 2-13

Addressing processes

- ❖ to receive messages, process must have *identifier*
- ❖ host device has unique 32-bit IP address
- ❖ *Q:* does IP address of host on which process runs suffice for identifying the process?
 - *A:* No, many processes can be running on same host
- ❖ *identifier* includes both IP address and port numbers associated with process on host.
- ❖ example port numbers:
 - HTTP server: 80
 - Mail server: 25
- ❖ to send HTTP message to gaia.cs.umass.edu web server:
 - IP address: 128.119.245.12
 - Port number: 80
- ❖ more shortly...

Application 2-14

App-layer protocol defines

- ❖ types of messages exchanged,
 - e.g., request, response
- ❖ message syntax:
 - what fields in messages & how fields are delineated
- ❖ message semantics
 - meaning of information in fields
- ❖ rules for when and how processes send & respond to messages
- public-domain protocols:**
 - ❖ defined in RFCs
 - ❖ allows for interoperability
 - ❖ e.g., HTTP, SMTP
- proprietary protocols:**
 - ❖ e.g., Skype

Application 2-15

What transport service does an app need?

Data loss

- ❖ some apps (e.g., audio) can tolerate some loss
- ❖ other apps (e.g., file transfer, telnet) require 100% reliable data transfer

Timing

- ❖ some apps (e.g., Internet telephony, interactive games) require low delay to be "effective"

Throughput

- ❖ some apps (e.g., multimedia) require minimum amount of throughput to be "effective"
- ❖ other apps ("elastic apps") make use of whatever throughput they get

Security

- ❖ encryption, data integrity, ...

Application 2-16

Transport service requirements of common apps

Application	Data loss	Throughput	Time Sensitive
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video: 10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

Application 2-17

Internet transport protocols services

TCP service:

- ❖ *connection-oriented*: setup required between client and server processes
- ❖ *reliable transport* between sending and receiving process
- ❖ *flow control*: sender won't overwhelm receiver
- ❖ *congestion control*: throttle sender when network overloaded
- ❖ *does not provide*: timing, minimum throughput guarantees, security

UDP service:

- ❖ unreliable data transfer between sending and receiving process
- ❖ does not provide: connection setup, reliability, flow control, congestion control, timing, throughput guarantee, or security

Q: why bother? Why is there a UDP?

Application 2-18

Internet apps: application, transport protocols

Application	Application layer protocol	Underlying transport protocol
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	HTTP (e.g., YouTube), RTP [RFC 1889]	TCP or UDP
Internet telephony	SIP, RTP, proprietary (e.g., Skype)	typically UDP

Application 2-19

Chapter 2: Application layer

2.1 Principles of network applications

- app architectures
- app requirements

2.2 Web and HTTP

2.3 FTP

2.4 Electronic Mail

- SMTP, POP3, IMAP

2.5 DNS

2.6 P2P applications

2.7 Socket programming with TCP

2.8 Socket programming with UDP

Application 2-20

Web and HTTP

First, a review...

- ❖ **web page** consists of **objects**
- ❖ object can be HTML file, JPEG image, Java applet, audio file,...
- ❖ web page consists of **base HTML-file** which includes several referenced objects
- ❖ each object is addressable by a **URL**
- ❖ example URL:


www.someschool.edu / someDept/pic.gif
host name path name

Application 2-21

HTTP overview

HTTP: hypertext transfer protocol

- ❖ Web's application layer protocol
- ❖ client/server model
 - **client**: browser that requests, receives, "displays" Web objects
 - **server**: Web server sends objects in response to requests



Application 2-22

HTTP overview (continued)

Uses TCP:

- ❖ client initiates TCP connection (creates socket) to server, port 80
- ❖ server accepts TCP connection from client
- ❖ HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server)
- ❖ TCP connection closed

HTTP is "stateless"

- ❖ server maintains no information about past client requests

aside
protocols that maintain "state" are complex!

- ❖ past history (state) must be maintained
- ❖ if server/client crashes, their views of "state" may be inconsistent, must be reconciled

Application 2-23

HTTP connections

non-persistent HTTP

- ❖ at most one object sent over TCP connection.

persistent HTTP

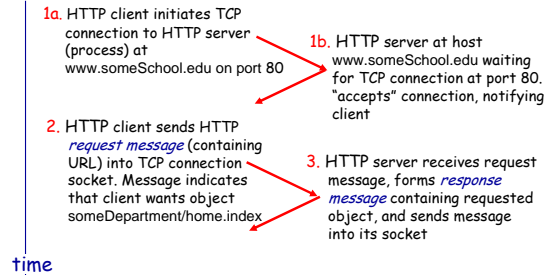
- ❖ multiple objects can be sent over single TCP connection between client, server.

Application 2-24

Nonpersistent HTTP

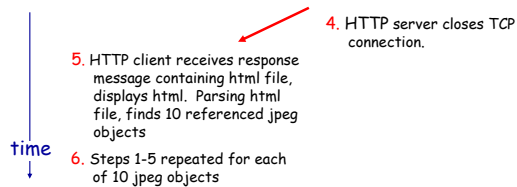
suppose user enters URL:

`www.someSchool.edu/someDepartment/home.index` (contains text, references to 10 jpeg images)



Application 2-25

Nonpersistent HTTP (cont.)



Application 2-26


Non-Persistent HTTP: Response time

definition of RTT: time for a small packet to travel from client to server and back.

response time:

- ❖ one RTT to initiate TCP connection
- ❖ one RTT for HTTP request and first few bytes of HTTP response to return
- ❖ file transmission time

total = $2RTT + \text{transmit time}$



Application 2-27

Persistent HTTP

non-persistent HTTP issues:

- ❖ requires 2 RTTs per object
- ❖ OS overhead for *each* TCP connection
- ❖ browsers often open parallel TCP connections to fetch referenced objects

persistent HTTP

- ❖ server leaves connection open after sending response
- ❖ subsequent HTTP messages between same client/server sent over open connection
- ❖ client sends requests as soon as it encounters a referenced object
- ❖ as little as one RTT for all the referenced objects

Application 2-28

HTTP request message

- ❖ two types of HTTP messages: *request, response*
- ❖ **HTTP request message:**
 - ASCII (human-readable format)

<initial line, different for request vs. response>

Header1: value1

Header2: value2

Header3: value3

<optional message body goes here, like file contents or query data;

it can be many lines long, or even binary data \$&*&@!^\$@>

Application 2-29

HTTP request

- ❖ Initial response line
GET /path/to/file/index.html HTTP/1.0
- ❖ The header name is not case-sensitive (the value maybe).
- ❖ Any number of spaces or tabs between : and value
- ❖ Header lines beginning with space or tab are actually part of the previous header line folded into multiple lines for easy reading.
- ❖ HTTP 1.0 defines 16 headers (non is required), while HTTP 1.1 defines 46 and one is required (Host:)

HTTP request message

❖ two types of HTTP messages: *request, response*

❖ **HTTP request message:**

- ASCII (human-readable format)

request line
(GET, POST,
HEAD commands)

header
lines


carriage return,
line feed at start
of line indicates
end of header lines

```
GET /index.html HTTP/1.1\r\n
Host: www-net.cs.umass.edu\r\n
User-Agent: Firefox/3.6.10\r\n
Accept: text/html,application/xhtml+xml\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7\r\n
Keep-Alive: 115\r\n
Connection: keep-alive\r\n
\r\n
```

carriage return character
line-feed character

Application 2-31

HTTP request message: general format



Application 2-32

HTTP 1.0 Examples

http://www.somehost.com/path/file.html

```
GET /path/file.html HTTP/1.0
From: someuser@jmarshall.com
User-Agent: HTTPTool/1.0
[blank line here]
```

```
HTTP/1.0 200 OK
Date: Fri, 31 Dec 1999 23:59:59 GMT
Content-Type: text/html
Content-Length: 1354
```

```
<html>
<body>
<h1>Happy New Millennium!</h1>
(more file contents)
.
.
.
</body>
</html>
```

Uploading form input

POST method:

- web page often includes form input
- ❖ input is uploaded to server in entity body

URL method:

- ❖ uses GET method
- ❖ input is uploaded in URL field of request
line: `www.somesite.com/animalsearch?monkeys&banana`

Application 2-34

Method types

HTTP/1.0

- ❖ GET
- ❖ POST
- ❖ HEAD
 - asks server to leave requested object out of response

HTTP/1.1

- ❖ GET, POST, HEAD
- ❖ PUT
 - uploads file in entity body to path specified in URL field
- ❖ DELETE
 - deletes file specified in the URL field

Application 2-35

HTTP response message

status line
(protocol
status code
status phrase)

header
lines

data, e.g.,
requested
HTML file

```
HTTP/1.1 200 OK\r\n
Date: Sun, 26 Sep 2010 20:09:20 GMT\r\n
Server: Apache/2.0.52 (CentOS)\r\n
Last-Modified: Tue, 30 Oct 2007 17:00:02 GMT\r\n
ETag: "17dc6-a5c-bf716880"\r\n
Accept-Ranges: bytes\r\n
Content-Length: 2652\r\n
Keep-Alive: timeout=10, max=100\r\n
Connection: Keep-Alive\r\n
Content-Type: text/html; charset=ISO-8859-1\r\n
\r\n
data data data data data ...
```

Application 2-36

HTTP response status codes

- ❖ status code appears in 1st line in server->client response message.
- ❖ some sample codes:
 - 200 OK**
 - request succeeded, requested object later in this msg
 - 301 Moved Permanently**
 - requested object moved, new location specified later in this msg (Location:)
 - 400 Bad Request**
 - request msg not understood by server
 - 404 Not Found**
 - requested document not found on this server
 - 505 HTTP Version Not Supported**

Application 2-37

HTTP 1.1

- ❖ In HTTP 1.1, one server with one IP address can be the home of several web domains.
- ❖ A request must specify which web domains it addresses.
- ❖ The header Host: must be included

Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

telnet cis.poly.edu 80 opens TCP connection to port 80 (default HTTP server port) at cis.poly.edu. anything typed in sent to port 80 at cis.poly.edu

2. type in a GET HTTP request:

GET /~ross/ HTTP/1.1
Host: cis.poly.edu by typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. look at response message sent by HTTP server!
(or use Wireshark!)

Application 2-39

Example

```
tigger 107 % telnet www.cse.yorku.ca 80
Trying 130.63.92.30...
Connected to www.cse.yorku.ca.
Escape character is '^['.
GET /course_archive/2011-12/W/3214/test.html HTTP/1.0 request

HTTP/1.1 200 OK
Date: Mon, 12 Dec 2011 17:57:18 GMT
Server: Apache/2.2.20 (Unix) DAV/2 mod_ssl/2.2.20 OpenSSL/0.9.8q PHP/5.2.17
X-Powered-By: PHP/5.2.17
Content-Length: 195
Connection: close
Content-Type: text/html response

<HTML>
<HEAD>
<TITLE>Archive of Web Pages</TITLE>
</HEAD>
<BODY>
<H2>
<CENTER>
<H2>Archive of Computer Science & Engineering Course Web Pages</H2>
</CENTER>
this is a simple text
</BODY>
</HTML>
Connection closed by foreign host.
tigger 108 %
```

User-server state: cookies

many Web sites use cookies

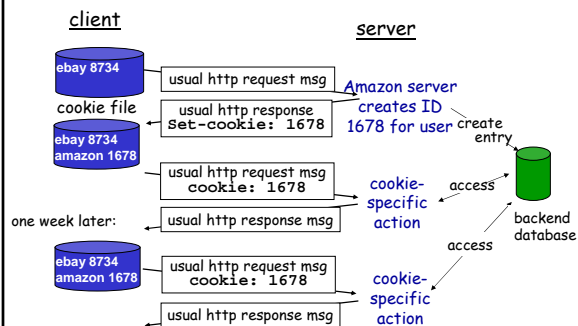
four components:

- 1) cookie header line of HTTP *response* message
- 2) cookie header line in HTTP *request* message
- 3) cookie file kept on user's host, managed by user's browser
- 4) back-end database at Web site

example:

- ❖ Susan always access Internet from PC
- ❖ visits specific e-commerce site for first time
- ❖ when initial HTTP requests arrives at site, site creates:
 - unique ID
 - entry in backend database for ID

Cookies: keeping "state" (cont.)



Cookies (continued)

what cookies can bring:

- ❖ authorization
- ❖ shopping carts
- ❖ recommendations
- ❖ user session state (Web e-mail)

how to keep "state":

- ❖ protocol endpoints: maintain state at sender/receiver over multiple transactions
- ❖ cookies: http messages carry state

cookies and privacy:


- ❖ cookies permit sites to learn a lot about you
- ❖ you may supply name and e-mail to sites

Application 2-43

Web caches (proxy server)

Goal: satisfy client request without involving origin server

- ❖ user sets browser: Web accesses via cache
- ❖ browser sends all HTTP requests to cache
 - object in cache: cache returns object
 - else cache requests object from origin server, then returns object to client



Application 2-44

More about Web caching

- ❖ cache acts as both client and server
- ❖ typically cache is installed by ISP (university, company, residential ISP)

why Web caching?

- ❖ reduce response time for client request
- ❖ reduce traffic on an institution's access link.
- ❖ Internet dense with caches: enables "poor" content providers to effectively deliver content (but so does P2P file sharing)

Application 2-45


Caching example

assumptions

- ❖ average object size = 100,000 bits
- ❖ avg. request rate from institution's browsers to origin servers = 15/sec
- ❖ delay from institutional router to any origin server and back to router = 2 sec

consequences

- ❖ utilization on LAN = 15%
- ❖ utilization on access link = 100%
- ❖ total delay = Internet delay + access delay + LAN delay
= 2 sec + minutes + milliseconds



Application 2-46


Caching example (cont)

possible solution

- ❖ increase bandwidth of access link to, say, 10 Mbps

consequence

- ❖ utilization on LAN = 15%
- ❖ utilization on access link = 15%
- ❖ Total delay = Internet delay + access delay + LAN delay
= 2 sec + msec + msec
- ❖ often a costly upgrade



Application 2-47


Caching example (cont)

possible solution:

- ❖ install cache

consequence


- ❖ suppose hit rate is 0.4
 - 40% requests will be satisfied almost immediately
 - 60% requests satisfied by origin server
- ❖ utilization of access link reduced to 60%, resulting in negligible delays (say 10 msec)
- ❖ total avg delay = Internet delay + access delay + LAN delay
= .6*(2.01) secs + .4*milliseconds < 1.4 secs



Application 2-48

Conditional GET

- ❖ **Goal:** don't send object if cache has up-to-date cached version
- ❖ **cache:** specify date of cached copy in HTTP request
If-modified-since: <date>
- ❖ **server:** response contains no object if cached copy is up-to-date:
HTTP/1.0 304 Not Modified




Application 2-49

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 **FTP**
- 2.4 Electronic mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P applications
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP

Application 2-50

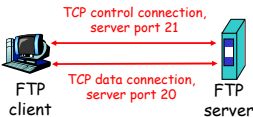
FTP: the file transfer protocol



- ❖ transfer file to/from remote host
- ❖ client/server model
 - **client:** side that initiates transfer (either to/from remote)
 - **server:** remote host
- ❖ ftp: RFC 959
- ❖ ftp server: port 21

Application 2-51

FTP: separate control, data connections

- ❖ FTP client contacts FTP server at port 21, TCP is transport protocol
 - ❖ client authorized over control connection
 - ❖ client browses remote directory by sending commands over control connection.
 - ❖ when server receives file transfer command, server opens 2nd TCP connection (for file) to client
 - ❖ after transferring one file, server closes data connection.
- 
- The diagram illustrates the FTP architecture. On the left is the 'FTP client' (represented by a laptop icon) and on the right is the 'FTP server' (represented by a server rack icon). Two red arrows connect them: the top arrow is labeled 'TCP control connection, server port 21' and the bottom arrow is labeled 'TCP data connection, server port 20'.
- ❖ server opens another TCP data connection to transfer another file.
 - ❖ control connection: "out of band"
 - ❖ FTP server maintains "state": current directory, earlier authentication

Application 2-52

FTP commands, responses

sample commands:

- ❖ sent as ASCII text over control channel
- ❖ USER *username*
- ❖ PASS *password*
- ❖ LIST return list of file in current directory
- ❖ RETR *filename* retrieves (gets) file
- ❖ STOR *filename* stores (puts) file onto remote host

sample return codes

- ❖ status code and phrase (as in HTTP)
- ❖ 331 Username OK, password required
- ❖ 125 data connection already open; transfer starting
- ❖ 425 Can't open data connection
- ❖ 452 Error writing file

Application 2-53

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P applications
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP

Application 2-54


Electronic Mail

Three major components:

- ❖ user agents
- ❖ mail servers
- ❖ simple mail transfer protocol: SMTP

User Agent

- ❖ a.k.a. "mail reader"
- ❖ composing, editing, reading mail messages
- ❖ e.g., Outlook, elm, Mozilla Thunderbird, iPhone mail client
- ❖ outgoing, incoming messages stored on server




Application 2-55

Electronic Mail: mail servers

Mail Servers

- ❖ **mailbox** contains incoming messages for user
- ❖ **message queue** of outgoing (to be sent) mail messages
- ❖ **SMTP protocol** between mail servers to send email messages
 - client: sending mail server
 - "server": receiving mail server



Application 2-56


Electronic Mail: SMTP [RFC 2821]

- ❖ uses TCP to reliably transfer email message from client to server, port 25
- ❖ direct transfer: sending server to receiving server
- ❖ three phases of transfer
 - handshaking (greeting)
 - transfer of messages
 - closure
- ❖ command/response interaction
 - **commands**: ASCII text
 - **response**: status code and phrase
- ❖ messages must be in 7-bit ASCII

Application 2-57

Scenario: Alice sends message to Bob

- 1) Alice uses UA to compose message and "to" bob@someschool.edu
- 2) Alice's UA sends message to her mail server; message placed in message queue
- 3) Client side of SMTP opens TCP connection with Bob's mail server
- 4) SMTP client sends Alice's message over the TCP connection
- 5) Bob's mail server places the message in Bob's mailbox
- 6) Bob invokes his user agent to read message



Application 2-58

Sample SMTP interaction

```
S: 220 hamburger.edu
C: HELO crepes.fr
S: 250 Hello crepes.fr, pleased to meet you
C: MAIL FROM: <alice@crepes.fr>
S: 250 alice@crepes.fr... Sender ok
C: RCPT TO: <bob@hamburger.edu>
S: 250 bob@hamburger.edu ... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Do you like ketchup?
C: How about pickles?
C: .
S: 250 Message accepted for delivery
C: QUIT
S: 221 hamburger.edu closing connection
```

Application 2-59

Try SMTP interaction for yourself:

- ❖ telnet servername 25
 - ❖ see 220 reply from server
 - ❖ enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands
- above lets you send email without using email client (reader)

Application 2-60

SMTP: final words


- ❖ SMTP uses persistent connections
 - ❖ SMTP requires message (header & body) to be in 7-bit ASCII
 - ❖ SMTP server uses CRLF, CRLF to determine end of message
- comparison with HTTP:**
- ❖ HTTP: pull
 - ❖ SMTP: push
 - ❖ both have ASCII command/response interaction, status codes
 - ❖ HTTP: each object encapsulated in its own response msg
 - ❖ SMTP: multiple objects sent in multipart msg

Application 2-61

Mail message format


SMTP: protocol for exchanging email msgs
RFC 822: standard for text message format:

- ❖ header lines, e.g.,
 - To:
 - From:
 - Subject:
- different from SMTP commands!*
- ❖ body
 - the "message", ASCII characters only



Application 2-62

Mail access protocols



- ❖ SMTP: delivery/storage to receiver's server
- ❖ mail access protocol: retrieval from server
 - POP: Post Office Protocol [RFC 1939]
 - authorization (agent <-->server) and download
 - IMAP: Internet Mail Access Protocol [RFC 1730]
 - more features (more complex)
 - manipulation of stored msgs on server
 - HTTP: gmail, Hotmail, Yahoo! Mail, etc.

Application 2-63

POP3 protocol

authorization phase

- ❖ client commands:
 - user: declare username
 - pass: password
- ❖ server responses
 - +OK
 - -ERR

transaction phase, client:

- ❖ list: list message numbers
- ❖ retr: retrieve message by number
- ❖ dele: delete
- ❖ quit

```
S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on

C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 1 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off
```

Application 2-64

POP3 (more) and IMAP

more about POP3

- ❖ previous example uses "download and delete" mode.
- ❖ Bob cannot re-read e-mail if he changes client
- ❖ "download-and-keep": copies of messages on different clients
- ❖ POP3 is stateless across sessions

IMAP

- ❖ keeps all messages in one place: at server
- ❖ allows user to organize messages in folders
- ❖ keeps user state across sessions:
 - names of folders and mappings between message IDs and folder name

Application 2-65

Chapter 2: Application layer

- ❖ 2.1 Principles of network applications
- ❖ 2.2 Web and HTTP
- ❖ 2.3 FTP
- ❖ 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- ❖ 2.5 DNS
- ❖ 2.6 P2P applications
- ❖ 2.7 Socket programming with TCP
- ❖ 2.8 Socket programming with UDP

Application 2-66

DNS: Domain Name System

people: many identifiers:

- SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) - used for addressing datagrams
- "name", e.g., www.yahoo.com - used by humans

Q: map between IP address and name, and vice versa ?

Domain Name System:

- ❖ *distributed database* implemented in hierarchy of many *name servers*
- ❖ *application-layer protocol* host, routers, name servers to communicate to *resolve* names (address/name translation)
 - note: core Internet function, implemented as application-layer protocol
 - complexity at network's "edge"

Application 2-67

DNS

DNS services

- ❖ hostname to IP address translation
- ❖ host aliasing
 - Canonical, alias names
- ❖ mail server aliasing
- ❖ load distribution
 - replicated Web servers: set of IP addresses for one canonical name


Why not centralize DNS?

- ❖ single point of failure
- ❖ traffic volume
- ❖ distant centralized database
- ❖ maintenance

doesn't *scale*!

Application 2-68

Distributed, Hierarchical Database



client wants IP for www.amazon.com; 1st approx:

- ❖ client queries a root server to find com DNS server
- ❖ client queries com DNS server to get amazon.com DNS server
- ❖ client queries amazon.com DNS server to get IP address for www.amazon.com

Application 2-69

DNS: Root name servers

- ❖ contacted by local name server that can not resolve name
- ❖ root name server:
 - contacts authoritative name server if name mapping not known
 - gets mapping
 - returns mapping to local name server



Application 2-70

TLD and Authoritative Servers

Top-level domain (TLD) servers:

- responsible for com, org, net, edu, aero, jobs, museums, and all top-level country domains, e.g.: uk, fr, ca, jp
- Network Solutions maintains servers for com TLD
- Educause for edu TLD

Authoritative DNS servers:

- organization's DNS servers, providing authoritative hostname to IP mappings for organization's servers (e.g., Web, mail).
- can be maintained by organization or service provider

Application 2-71

Local Name Server

- ❖ does not strictly belong to hierarchy
- ❖ each ISP (residential ISP, company, university) has one
 - also called "default name server"
- ❖ when host makes DNS query, query is sent to its local DNS server
 - acts as proxy, forwards query into hierarchy


Application 2-72

DNS name resolution example

- ❖ host at cis.poly.edu wants IP address for gaia.cs.umass.edu

iterated query:

- ❖ contacted server replies with name of server to contact
- ❖ "I don't know this name, but ask this server"

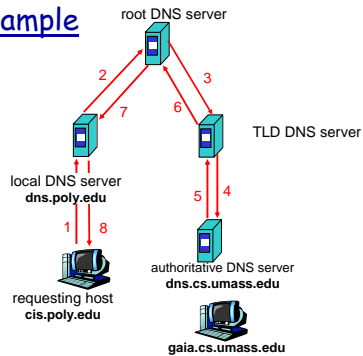


Application 2-73

DNS name resolution example

recursive query:

- ❖ puts burden of name resolution on contacted name server
- ❖ heavy load?



Application 2-74

DNS: caching and updating records

- ❖ once (any) name server learns mapping, it *caches* mapping
 - cache entries timeout (disappear) after some time
 - TLD servers typically cached in local name servers
 - Thus root name servers not often visited
- ❖ update/notify mechanisms proposed IETF standard
 - RFC 2136

Application 2-75

DNS records

DNS: distributed db storing resource records (RR)

RR format: (name, value, type, ttl)

Type=A

- name is hostname
- value is IP address

Type=NS

- name is domain (e.g., foo.com)
- value is hostname of authoritative name server for this domain

Type=CNAME

- name is alias name for some "canonical" (the real) name
- www.ibm.com is really servereast.backup2.ibm.com
- value is canonical name

Type=MX

- value is name of mailserver associated with name

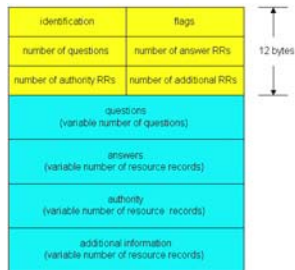
Application 2-76

DNS protocol, messages

DNS protocol: *query* and *reply* messages, both with same *message format*

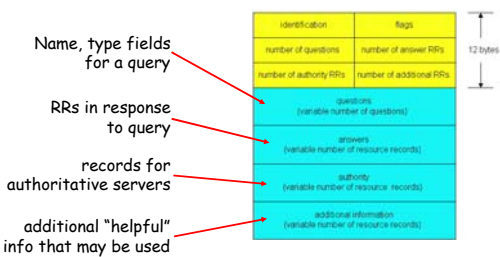
msg header

- ❖ **identification**: 16 bit # for query, reply to query uses same #
- ❖ **flags**:
 - query or reply
 - recursion desired
 - recursion available
 - reply is authoritative



Application 2-77

DNS protocol, messages



Application 2-78

Inserting records into DNS

- ❖ example: new startup "Network Utopia"
- ❖ register name networkutopia.com at *DNS registrar* (e.g., Network Solutions)
 - provide names, IP addresses of authoritative name server (primary and secondary)
 - registrar inserts two RRs into com TLD server:

```
(networkutopia.com, dns1.networkutopia.com, NS)
(dns1.networkutopia.com, 212.212.212.1, A)
```

- ❖ create authoritative server Type A record for www.networkutopia.com; Type MX record for networkutopia.com
- ❖ *How do people get IP address of your Web site?*

Application 2-79

Chapter 2: Application layer

- | | |
|--|---------------------------------|
| 2.1 Principles of network applications | 2.6 P2P applications |
| 2.2 Web and HTTP | 2.7 Socket programming with TCP |
| 2.3 FTP | 2.8 Socket programming with UDP |
| 2.4 Electronic Mail <ul style="list-style-type: none">▪ SMTP, POP3, IMAP | |
| 2.5 DNS | |


Application 2-80

Pure P2P architecture

- ❖ *no* always-on server
- ❖ arbitrary end systems directly communicate
- ❖ peers are intermittently connected and change IP addresses

Three topics:


- file distribution
- searching for information
- case Study: Skype



Application 2-81

File Distribution: Server-Client vs P2P


Question: How much time to distribute file from one server to N peers?



Application 2-82

File distribution time: server-client

- ❖ server sequentially sends N copies:
 - NF/u_s time
- ❖ client i takes F/d_i time to download




Time to distribute F to N clients using client/server approach

increases linearly in N (for large N)

Application 2-83

File distribution time: P2P

- ❖ server must send one copy: F/u_s time
- ❖ client i takes F/d_i time to download
- ❖ NF bits must be downloaded (aggregate)
 - fastest possible upload rate: $u_s + \sum u_i$




$$d_{P2P} = \max \{ F/u_s, F/\min(d_i), NF/(u_s + \sum u_i) \}$$

Application 2-84

Server-client vs. P2P: example

Client upload rate = u , $F/u = 1$ hour, $u_s = 10u$, $d_{\min} \geq u_s$




Application 2-85

File distribution: BitTorrent

P2P file distribution


tracker: tracks peers participating in torrent

torrent: group of peers exchanging chunks of a file



Application 2-86

BitTorrent (1)



- ❖ file divided into 256KB *chunks*.
- ❖ peer joining torrent:
 - has no chunks, but will accumulate them over time
 - registers with tracker to get list of peers, connects to subset of peers ("neighbors")
- ❖ while downloading, peer uploads chunks to other peers.
- ❖ peers may come and go
- ❖ once peer has entire file, it may (selfishly) leave or (altruistically) remain

Application 2-87

BitTorrent (2)

Pulling Chunks

- ❖ at any given time, different peers have different subsets of file chunks
- ❖ periodically, a peer (Alice) asks each neighbor for list of chunks that they have.
- ❖ Alice sends requests for her missing chunks
 - rarest first


Sending Chunks: tit-for-tat

- ❖ Alice sends chunks to four neighbors currently sending her chunks *at the highest rate*
 - re-evaluate top 4 every 10 secs
- ❖ every 30 secs: randomly select another peer, starts sending chunks
 - newly chosen peer may join top 4
 - "optimistically unchoke"

Application 2-88

BitTorrent: Tit-for-tat

- (1) Alice "optimistically unchokes" Bob
- (2) Alice becomes one of Bob's top-four providers; Bob reciprocates
- (3) Bob becomes one of Alice's top-four providers



Application 2-89

Distributed Hash Table (DHT)

- ❖ DHT: distributed P2P database
- ❖ database has (key, value) pairs;
 - key: ss number; value: human name
 - key: content type; value: IP address
- ❖ peers **query** DB with key
 - DB returns values that match the key
- ❖ peers can also **insert** (key, value) peers

Application 2-90

DHT Identifiers

- ❖ assign integer identifier to each peer in range $[0, 2^n - 1]$.
 - Each identifier can be represented by n bits.
- ❖ require each key to be an integer in **same range**.
- ❖ to get integer keys, hash original key.
 - e.g., key = $h(\text{"Led Zeppelin IV"})$
 - this is why they call it a distributed "hash" table

Application 2-91

How to assign keys to peers?


- ❖ central issue:
 - assigning (key, value) pairs to peers.
- ❖ rule: assign key to the peer that has the **closest** ID.
- ❖ convention in lecture: closest is the **immediate successor** of the key.
- ❖ e.g., $n=4$; peers: 1,3,4,5,8,10,12,14;
 - key = 13, then successor peer = 14
 - key = 15, then successor peer = 1

Application 2-92

Adding a key

- ❖ Assume that a peer wants to insert a record in the database.
- ❖ Simply calculate the hash, and send it to the immediate successor.
- ❖ How can we know the immediate successor?
- ❖ Every peer keeps track of all the peers is not a viable solution (may be in the millions).

Circular DHT (1)




- ❖ each peer *only* aware of immediate successor and predecessor.
- ❖ "overlay network"

Application 2-94

Circular DHT (2)


$O(N)$ messages on avg to resolve query, when there are N peers

Define closest as closest successor



Application 2-95


Circular DHT with Shortcuts



- ❖ each peer keeps track of IP addresses of predecessor, successor, short cuts.
- ❖ reduced from 6 to 2 messages.
- ❖ possible to design shortcuts so $O(\log N)$ neighbors, $O(\log N)$ messages in query

Application 2-96

Peer Churn




- ❖ To handle peer churn, require each peer to know the IP address of its two successors.
- ❖ Each peer periodically pings its two successors to see if they are still alive.

- ❖ peer 5 abruptly leaves
- ❖ Peer 4 detects; makes 8 its immediate successor; asks 8 who its immediate successor is; makes 8's immediate successor its second successor.
- ❖ What if peer 13 wants to join? *Only knows of node 1*

Application 2-97

P2P Case study: Skype


- ❖ inherently P2P: pairs of users communicate.
- ❖ proprietary application-layer protocol (inferred via reverse engineering)- All messages encrypted
- ❖ hierarchical overlay with SNs
- ❖ Index maps usernames to IP addresses; distributed over SNs
- ❖ To call someone, search the distributed index for his/her IP



Application 2-98

Peers as relays

- ❖ problem when both Alice and Bob are behind "NATs".
 - NAT prevents an outside peer from initiating a call to insider peer
- ❖ solution:
 - When you login, you are assigned a nonNAT-ed SN
 - using Alice's and Bob's SNs, *relay* is chosen
 - each peer initiates session with relay.
 - peers can now communicate through NATs via relay



Application 2-99

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P applications
- 2.7 **Socket programming with TCP**
- 2.8 Socket programming with UDP

Application 2-100

Socket programming

Goal: learn how to build client/server application that communicate using sockets

Socket API

- ❖ introduced in BSD4.1 UNIX, 1981
- ❖ explicitly created, used, released by apps
- ❖ client/server paradigm
- ❖ two types of transport service via socket API:
 - unreliable datagram
 - reliable, byte stream-oriented

socket


a *host-local, application-created, OS-controlled* interface (a "door") into which application process can **both send and receive** messages to/from another application process

Application 2-101

Socket-programming using TCP

Socket: a door between application process and end-end-transport protocol (TCP or UDP)

TCP service: reliable transfer of *bytes* from one process to another



Application 2-102

Socket programming *with TCP*

Client must contact server

- ❖ server process must first be running
- ❖ server must have created socket (door) that welcomes client's contact

Client contacts server by:

- ❖ creating client-local TCP socket
- ❖ specifying IP address, port number of server process
- ❖ when **client creates socket**: client TCP establishes connection to server TCP

- ❖ when contacted by client, **server TCP creates new socket** for server process to communicate with client
 - allows server to talk with multiple clients
 - source port numbers used to distinguish clients (more in Chap 3)

application viewpoint

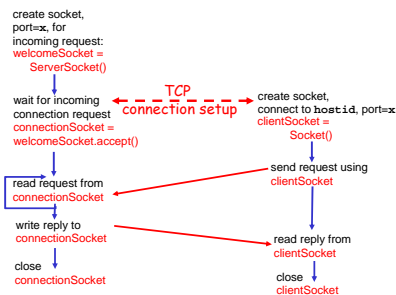
TCP provides reliable, in-order transfer of bytes ("pipe") between client and server

Application 2-103

Client/server socket interaction: TCP

Server (running on `hostid`)


Client



Application 2-104

Stream jargon

- ❖ **stream** is a sequence of characters that flow into or out of a process.
- ❖ **input stream** is attached to some input source for the process, e.g., keyboard or socket.
- ❖ **output stream** is attached to an output source, e.g., monitor or socket.



Application 2-105

Socket programming with TCP

Example client-server app:

- 1) client reads line from standard input (`inFromUser` stream), sends to server via socket (`outToServer` stream)
- 2) server reads line from socket
- 3) server converts line to uppercase, sends back to client
- 4) client reads, prints modified line from socket (`inFromServer` stream)

Application 2-106

Example: Java client (TCP)

```
import java.io.*;
import java.net.*;
class TCPCClient {

    public static void main(String argv[]) throws Exception
    {
        String sentence;
        String modifiedSentence;

        // create input stream
        BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));

        // create clientSocket object of type Socket, connect to server
        Socket clientSocket = new Socket("hostname", 6789);

        // create output stream attached to socket
        DataOutputStream outToServer =
            new DataOutputStream(clientSocket.getOutputStream());
```

← This package defines `Socket()` and `ServerSocket()` classes

server name,
e.g., `www.umass.edu`

server port #

create input stream
create clientSocket object of type `Socket`, connect to server
create output stream attached to socket

Application 2-107

Example: Java client (TCP), cont.

```
        // create input stream attached to socket
        BufferedReader inFromServer =
            new BufferedReader(new InputStreamReader(clientSocket.getInputStream()));

        sentence = inFromUser.readLine();

        // send line to server
        outToServer.writeBytes(sentence + '\n');

        // read line from server
        modifiedSentence = inFromServer.readLine();

        System.out.println("FROM SERVER: " + modifiedSentence);

        // close socket (clean up behind yourself)
        clientSocket.close();
    }
}
```

Application 2-108

Example: Java server (TCP)

```
import java.io.*;
import java.net.*;

class TCPServer {

    public static void main(String argv[]) throws Exception
    {
        String clientSentence;
        String capitalizedSentence;

        create
        welcoming socket
        at port 6789 → ServerSocket welcomeSocket = new ServerSocket(6789);

        wait, on welcoming
        socket accept() method
        for client contact create,
        new socket on return → while(true) {
                                Socket connectionSocket = welcomeSocket.accept();

                                create input
                                stream, attached
                                to socket → BufferedReader inFromClient =
                                        new BufferedReader(new
                                        InputStreamReader(connectionSocket.getInputStream()));
```

Application 2-109

Example: Java server (TCP), cont

```
        create output
        stream, attached
        to socket → DataOutputStream outToClient =
                    new DataOutputStream(connectionSocket.getOutputStream());

        read in line
        from socket → clientSentence = inFromClient.readLine();

        capitalizedSentence = clientSentence.toUpperCase() + '\n';

        write out line
        to socket → outToClient.writeBytes(capitalizedSentence);
    }
} → end of while loop,
    loop back and wait for
    another client connection
```

Application 2-110

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P applications
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP

Application 2-111

Socket programming *with UDP*

UDP: no "connection" between client and server

- ❖ no handshaking
- ❖ sender explicitly attaches IP address and port of destination to each packet
- ❖ server must extract IP address, port of sender from received packet

UDP: transmitted data may be received out of order, or lost

application viewpoint:

UDP provides *unreliable* transfer of groups of bytes ("datagrams") between client and server

Application 2-112

Client/server socket interaction: UDP

Server (running on `hostid`)

Client

create socket,
port= x,
`serverSocket = DatagramSocket()`

read datagram from
`serverSocket`

write reply to
`serverSocket`
specifying
client address,
port number

create socket,
`clientSocket = DatagramSocket()`


Create datagram with server IP and
port=x; send datagram via
`clientSocket`

read datagram from
`clientSocket`

close
`clientSocket`

Application 2-113

Example: Java client (UDP)



Application 2-114

Example: Java client (UDP)

```
import java.io.*;
import java.net.*;

class UDPClient {
    public static void main(String args[]) throws Exception
    {
        create input stream → BufferedReader inFromUser =
                                new BufferedReader(new InputStreamReader(System.in));
        create client socket → DatagramSocket clientSocket = new DatagramSocket();
        translate hostname to IP address using DNS → InetAddress IPAddress = InetAddress.getByName("hostname");

        byte[] sendData = new byte[1024];
        byte[] receiveData = new byte[1024];

        String sentence = inFromUser.readLine();
        sendData = sentence.getBytes();
    }
}
```

Application 2-115

Example: Java client (UDP), cont.

```
create datagram with data-to-send, length, IP addr, port → DatagramPacket sendPacket =
                                                                new DatagramPacket(sendData, sendData.length, IPAddress, 9876);
send datagram to server → clientSocket.send(sendPacket);
read datagram from server → DatagramPacket receivePacket =
                             new DatagramPacket(receiveData, receiveData.length);
                             clientSocket.receive(receivePacket);

String modifiedSentence =
    new String(receivePacket.getData());

System.out.println("FROM SERVER:" + modifiedSentence);
clientSocket.close();
}
```

Application 2-116

Example: Java server (UDP)

```
import java.io.*;
import java.net.*;

class UDPServer {
    public static void main(String args[]) throws Exception
    {
        create datagram socket at port 9876 → DatagramSocket serverSocket = new DatagramSocket(9876);

        byte[] receiveData = new byte[1024];
        byte[] sendData = new byte[1024];

        while(true)
        {
            create space for received datagram → DatagramPacket receivePacket =
                                                  new DatagramPacket(receiveData, receiveData.length);
            receive datagram → serverSocket.receive(receivePacket);
        }
    }
}
```

Application 2-117

Example: Java server (UDP), cont

```
String sentence = new String(receivePacket.getData());  
get IP addr } InetAddress IPAddress = receivePacket.getAddress();  
port #, of }  
sender } int port = receivePacket.getPort();  
  
String capitalizedSentence = sentence.toUpperCase();  
  
sendData = capitalizedSentence.getBytes();  
  
create datagram } DatagramPacket sendPacket =  
to send to client } new DatagramPacket(sendData, sendData.length, IPAddress,  
port);  
  
write out } serverSocket.send(sendPacket);  
datagram }  
to socket }  
}  
}  
end of while loop,  
loop back and wait for  
another datagram
```

Application 2-118

Network programming in C

- ❖ A very excellent tutorial is Beej's guide to network programming (see course web site)
- ❖ Here, I will present very minimal information just enough to write one server/client application.
- ❖ The above tutorial covers both IPv4 and IPv6, here I will cover only IPv4

Byte Order

- ❖ Big endian 0xb34f are represented as b3 in one byte, the next one contain 4f That also is network byte order
- ❖ Little endian 0xb34f are represented as 4f followed by b3 (x86 compatible machines)
- ❖ To prevent confusion, convert every thing before you send to network order and convert every thing that you receive to host order
- ❖ htons() htonl(), ntohs(), ntohl()

Structs

```
struct addrinfo {
    int         ai_flags;        // AI_PASSIVE, AI_CANONNAME, etc.
    int         ai_family;       // AF_INET, AF_INET6, AF_UNSPEC
    int         ai_socktype;     // SOCK_STREAM, SOCK_DGRAM
    int         ai_protocol;     // use 0 for "any"
    size_t      ai_addrlen;      // size of ai_addr in bytes
    struct sockaddr *ai_addr;    // struct sockaddr_in or _in6
    char        *ai_canonname;   // full canonical hostname

    struct addrinfo *ai_next;    // linked list, next node
};
```

- ❑ This is a new struct used to hold information needed by the socket.
- ❑ `getaddrinfo()` is used to fill it up

Structs

```
struct sockaddr {
    unsigned short sa_family;    // address family, AF_XXX
    char           sa_data[14];  // 14 bytes of protocol address
};
```

```
//IPv4 only
struct sockaddr_in {
    short int      sin_family;   // Address family, AF_INET
    unsigned short int sin_port; // Port number
    struct in_addr sin_addr;     // Internet address
    unsigned char  sin_zero[8]; // Same size as struct sockaddr
};
```

- ❖ Can be casted to each other

Structs

- ❖ Struct `sockaddr_storage` is large enough to hold both IPv4 and IPv6 info.
- ❖ Check the `ss_family`, then cast it to `sockaddr_in` or `sockaddr_in6`

```
struct sockaddr_storage {
    sa_family_t ss_family;    // address family

    // all this is padding, implementation specific, ignore it:
    char        __ss_pad1[__SS_PAD1SIZE];
    int64_t     __ss_align;
    char        __ss_pad2[__SS_PAD2SIZE];
};
```

Example

- ❖ A minimal code - no error checking
- ❖ The complete code is in the Beej's tutorial and is available at the course web site

Example

```
/* no error checking, only for IPv4 see web site for the full code */
#include <stdio.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netdb.h>
#include <arpa/inet.h>

int main(int argc, char *argv[])
{
    struct addrinfo hints, *res, *p;
    int status;
    char ipstr[INET6_ADDRSTRLEN];

    memset(&hints, 0, sizeof hints);
    hints.ai_family = AF_INET; // AF_INET for IPv4 only
    hints.ai_socktype = SOCK_STREAM;

    getaddrinfo(argv[1], NULL, &hints, &res);

    printf("IP addresses for %s:\n\n", argv[1]);
```

Example

```
for(p = res; p != NULL; p = p->ai_next) {
    void *addr;
    char *ipver;

    // get the pointer to the address itself
    struct sockaddr_in *ipv4 = (struct sockaddr_in *)p->ai_addr;
    addr = &(ipv4->sin_addr);
    ipver = "IPv4";

    // convert the IP to a string and print it:
    inet_ntop(p->ai_family, addr, ipstr, sizeof ipstr);

    printf(" %s: %s\n", ipver, ipstr);
}

freeaddrinfo(res); // free the linked list

return 0;
}
```

Example

tigger 121% gcc showipv4.c -lnsl
tigger 122% a.out indigo.cse.yorku.ca
IP addresses for indigo.cse.yorku.ca:

IPv4: 130.63.92.157
tigger 123% a.out www.cnn.com
IP addresses for www.cnn.com:

IPv4: 157.166.226.26
IPv4: 157.166.255.18
IPv4: 157.166.255.19
IPv4: 157.166.226.25

Client Server Example in C

```
/* client.c -- a stream socket client demo*/  
#include <stdio.h>  
#include <stdlib.h>  
#include <unistd.h>  
#include <errno.h>  
#include <string.h>  
#include <netdb.h>  
#include <sys/types.h>  
#include <netinet/in.h>  
#include <sys/socket.h>  
#include <arpa/inet.h>  
#define PORT "3490" // the port client will be connecting to  
#define MAXDATASIZE 100 // max number of bytes we can get at once  
// get sockaddr, IPv4 or IPv6:  
void *get_in_addr(struct sockaddr *sa) {  
    if (sa->sa_family == AF_INET) {  
        return &(((struct sockaddr_in*)sa)->sin_addr);  
    }  
    return &(((struct sockaddr_in6*)sa)->sin6_addr);  
}
```

Client Server cont.

```
int main(int argc, char *argv[])  
{  
    int sockfd, numbytes;  
    char buf[MAXDATASIZE];  
    struct addrinfo hints, *servinfo, *p;  
    int rv;  
    char s[INET6_ADDRSTRLEN];  
    if (argc != 2) {  
        fprintf(stderr, "usage: client hostname\n");  
        exit(1);  
    }  
    memset(&hints, 0, sizeof hints);  
    hints.ai_family = AF_UNSPEC;  
    hints.ai_socktype = SOCK_STREAM;  
  
    if ((rv = getaddrinfo(argv[1], PORT, &hints, &servinfo)) != 0) {  
        fprintf(stderr, "getaddrinfo: %s\n", gai_strerror(rv));  
        return 1;  
    }  
}
```

Client Server cont.

```
//Loop thorough all the results and connect to the first one
for(p = servinfo; p != NULL; p = p->ai_next) {
    if ((sockfd = socket(p->ai_family, p->ai_socktype,
        p->ai_protocol)) == -1) {
        perror("client: socket");
        continue;
    }

    if (connect(sockfd, p->ai_addr, p->ai_addrlen) == -1) {
        close(sockfd);
        perror("client: connect");
        continue;
    }
    break;
}
```

Client Server cont.

```
if (p == NULL) {    fprintf(stderr, "client: failed to connect\n");    return 2; }

inet_ntop(p->ai_family, get_in_addr((struct sockaddr *)p->ai_addr),
    s, sizeof s);
printf("client: connecting to %s\n", s);

freeaddrinfo(servinfo); // all done with this structure

if ((numbytes = recv(sockfd, buf, MAXDATASIZE-1, 0)) == -1) {
    perror("recv");
    exit(1);
}

buf[numbytes] = '\0';
printf("client: received %s\n", buf);
close(sockfd);
return 0;
}
```

UDP in C

❖ See the programs on the course web site.

Chapter 2: Summary

our study of network apps now complete!

- ❖ application architectures
 - client-server
 - P2P
 - hybrid
- ❖ application service requirements:
 - reliability, bandwidth, delay
- ❖ Internet transport service model
 - connection-oriented, reliable: TCP
 - unreliable, datagrams: UDP
- ❖ specific protocols:
 - HTTP
 - FTP
 - SMTP, POP, IMAP
 - DNS
 - P2P: BitTorrent, Skype
- ❖ socket programming

Application 2-133

Chapter 2: Summary

most importantly: learned about *protocols*

- ❖ typical request/reply message exchange:
 - client requests info or service
 - server responds with data, status code
- ❖ message formats:
 - headers: fields giving info about data
 - data: info being communicated
- Important themes:*
 - ❖ control vs. data msgs
 - ❖ in-band, out-of-band
 - ❖ centralized vs. decentralized
 - ❖ stateless vs. stateful
 - ❖ reliable vs. unreliable msg transfer
 - ❖ "complexity at network edge"

Application 2-134
